



Regular Article

Irrigation infrastructure and satellite-measured land cultivation impacts: Evidence from the Senegal river valley[☆]

Abdoulaye Cisse^a, Alain de Janvry^b, Joel Ferguson^c, Marco Gonzalez-Navarro^b ^{*,}
Samba Mbaye^d, Elisabeth Sadoulet^b, Mame Mor Anta Syll^d

^a The World Bank, Washington, DC, United States of America

^b University of California, Berkeley, United States of America

^c University of Wisconsin - Madison, United States of America

^d Université Gaston Berger, Saint-Louis, Senegal

ARTICLE INFO

Dataset link: <https://osf.io/2z6n4/>

JEL classification:

O13

O18

Q15

H54

Keywords:

Irrigation

Satellite images

Land use

Agriculture

Africa

Senegal

ABSTRACT

Expanding irrigation in sub-Saharan Africa is widely viewed as a promising strategy for closing yield gaps and enhancing resilience to climate change. Drawing on more than 3,000 satellite images over a 30-year period, we examine the impact of irrigation infrastructure development in the Senegal River Valley. We find that cultivation rates increase substantially following irrigation project completion. Cultivation rates are remarkably stable at around 25 percentage points above pre-irrigation levels for the first 20 years, and trend even higher from years 20 to 25. Moreover, we show that crops cultivated on irrigated land are significantly less sensitive to both positive and negative temperature shocks, underscoring the role of irrigation in climate adaptation. Despite these aggregate gains, we document considerable heterogeneity in project outcomes, with intermittent land use remaining widespread. To shed light on these patterns, we complement the satellite analysis with farmer survey data, which point to persistent water access constraints as a key barrier to continuous cultivation—constraints that cannot be resolved solely through individual farmer action.

1. Introduction

Food security in sub-Saharan Africa remains heavily reliant on an agricultural sector that continues to struggle to achieve high productivity, both in terms of yield levels and yield stability. Enhancing agricultural productivity is widely recognized as a critical priority, particularly in light of rapidly growing demand driven by population growth, which is projected to reach 3.3 billion in sub-Saharan Africa by the end of the century (UN, 2024). At the same time, climate change poses an additional threat to agricultural output and food security in the region (Lobell et al., 2008). In this context, a key strategy for addressing food deficits involves transitioning away from dependence on rainfed and flood-recession agriculture towards more widespread adoption of irrigated agriculture (Rosa et al., 2020).

In this paper, we document the consequences of the damming of the Senegal river in the Senegal River Valley in the 1980s on the expansion of cultivated land in Senegal using a long time-series of satellite imagery. Senegal provides a context that is uniquely amenable to analysis. First, the Senegal River Valley is the site of one of the few existing large irrigation schemes in sub-Saharan Africa, though more such projects have been proposed recently (Higginbottom et al., 2021). Second, the Sahel, through which the Senegal River runs, is a region that has grappled with desertification since the 1970s (Hein and De Ridder, 2006), and so may provide a useful example for studying the role of irrigation in climate change adaptation. The region's sparse background vegetation further facilitates the detection of cultivated land via satellite imagery. Finally, the development of the valley within

[☆] The authors gratefully acknowledge financial support from the BASIS Feed the Future Innovation Lab for Markets, Risk and Resilience through the United States Agency for International Development (USAID) grant number A20-1825-S003. IRB approval was obtained from North Star Review Board under protocol NB300149. The contents are the responsibility of the authors and do not necessarily reflect the views of USAID or the US Government. We thank the Senegalese land development authority (SAED) for sharing their administrative data on irrigation projects and Khaly Fall and Ousmane Samaké from the land development authority for insightful information and discussions on the results. Authors are listed in alphabetical order.

* Corresponding author.

E-mail addresses: acisse14@worldbank.org (A. Cisse), alain@berkeley.edu (A. de Janvry), ferguson7@wisc.edu (J. Ferguson), marcog@berkeley.edu (M. Gonzalez-Navarro), samba.mbaye@ugb.edu.sn (S. Mbaye), esadoulet@berkeley.edu (E. Sadoulet), syll.mame-mor-anta1@ugb.edu.sn (M.M.A. Syll).

<https://doi.org/10.1016/j.jdevec.2025.103615>

Received 31 August 2024; Received in revised form 8 August 2025; Accepted 15 August 2025

Available online 4 September 2025

0304-3878/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Senegal has been managed by a single government agency, which has systematically documented irrigation infrastructure investments since at least the 1970s, providing a rare and valuable source of administrative data.

In spite of its policy and academic interest, evaluating the role of irrigation infrastructure in most developing countries has proven challenging due to a lack of time series data on land use and agricultural output. Indeed, there are no high-resolution agricultural output or land use data for the Senegal River Valley spanning the decades over which irrigation development has occurred. To fill this data gap, we turn to a collection of over 3000 publicly available Landsat images collected between 1985 and 2019. To convert the imagery into useful data for studying agricultural outcomes, we calculate the normalized difference vegetation index (NDVI) of each 30 m × 30 m pixel. A large body of remote sensing research documents the usefulness of NDVI in detecting biomass, with greater values indicative of denser and healthier plants. We take the maximum NDVI level over either the rainy season or the non-rainy season (also called ‘off-season’), and apply a simple threshold rule to infer cultivation. We validate this procedure by documenting that, in areas lacking irrigation, NDVI very rarely reaches above our chosen threshold.

We combine our newly developed dataset on cultivated land with administrative data from the Senegalese government covering 860 distinct irrigation infrastructure projects completed between 1988 and 2019. These administrative records include information on the geographic boundaries of each project and its constituent subsections – referred to as “development units” – which are operated by different producer organizations or individual farmers.¹ The data also record the year of project completion and key project characteristics, such as whether the unit is managed privately or by the public land development authority, and whether it is classified as an agribusiness. Merging these sources yields a panel dataset at the development unit level, tracking seasonal cultivation rates over more than three decades and covering 1364 distinct units.

Because we observe cultivation activity before and after irrigation project completion, we can estimate the effect of such projects using an event study design. There are two methodological challenges that must be accounted for. First, because irrigation projects are completed in different years, treatment timing is staggered, leading to potential issues with the traditional two-way fixed effects estimator in the presence of unrestricted treatment effect heterogeneity (Sun and Abraham, 2021). Second, the earlier half of the study period has a relatively large share of missing data due to lack of high-quality imagery, making our dataset an unbalanced panel.

We overcome both of these issues using the imputation event study estimator developed by Borusyak et al. (2024), which provides us with consistent estimates of a weighted average treatment effect on the treated under general treatment effect heterogeneity in panels which may be unbalanced. We posit a simple two-way fixed effect specification of agricultural land utilization in the absence of irrigation infrastructure and show that when this model is augmented with pre-event indicators there is no evidence of pre-project trends, supporting the parallel trends assumption necessary for causal identification. Once the model of utilization in the absence of irrigation is estimated, treatment effect estimates are calculated by predicting untreated potential outcomes, subtracting them from observed treated potential outcomes, and averaging these differences across observations.

Our results provide strong evidence that cultivation rates increase substantially beginning in the first year after irrigation project completion. Beyond the first year, we find cultivation rates are remarkably stable at around 25 percentage points above pre-irrigation levels for the first 20 years, and trend even higher from years 20 to 25, especially

during the dry off-season. Considering that average cultivation rates prior to project completion were only 10 percent in the rainy season and 4 percent in the off-season across the sample, these are large and significant effects.

While our main results provide a positive picture of land cultivation patterns following the introduction of irrigation infrastructure, this large estimated average impact masks significant heterogeneity across development units and over time. Irrigation has generally had an increasing impact on cultivation since the year 2000, with an acceleration since 2010. Moreover, the management type (public versus private) and the unit’s designation as an agribusiness constitute significant determinants of observed cultivation rates.

We proceed to classify development units according to their cultivation performance post project completion. Exploiting our spatially and temporally granular dataset, for each development unit we calculate the share of pixels that change cultivation status (from zero to one and vice versa) between year t and $t-1$. We classify development units as having “intermittent land use” if we observe cultivation rates that vary substantially across years (more than 10 percent annual average status change). The remaining development units are classified into “Low Use” (with an average cultivation rate below 33 percent), “Medium Use” (cultivation rate between 33 and 67 percent), and “High Use” (cultivation rate above 67 percent).

Using this classification, we observe that 18 percent of development units are Low Use and 36 percent Intermittent Use, so that more than half of the land opened to irrigation is underutilized in some form. Medium Use account for 31 percent, and only 14 percent of development units have High Use cultivation rates. This heterogeneity in cultivation rates implies the existence of a large inefficiency in land use, as we do not find evidence that intermittent use is due to land rotation.

To explore the causes for such widespread under-utilization of irrigated land, we move from satellite data to our dedicated farmer survey. Our sample consists of farmers that were known in their community for having some land that was either intermittently used or continuously unused. We obtained information from farmers about the history of land use in their plots during the three off-seasons in the 2021–2023 period as well as the causes for non-use.

We find that, in spite of all plots in our sample being in projects with completed irrigation infrastructure, lack of water availability was cited as the main constraint to land cultivation. This is the case for 61.6 percent of continuously unused land and 34.7 percent of intermittently used land across these three seasons. The second most important cause for lack of utilization is financial constraints, especially among intermittently used plots.

Further analysis shows that the water access problem causing non-use is not seen as something individual farmers can address on their own. In almost all cases, the solution is seen as requiring the intervention of actors external to the development unit, such as those insuring the maintenance of major canals that serve multiple development units.

Our study offers a proof of concept for integrating satellite-based remote sensing analyses with targeted field surveys to deepen our understanding of irrigation infrastructure impacts. In doing so, we confirm that a substantial share of irrigable land remains either persistently unused or subject to intermittent use. Moreover, farm-level interviews reveal that limited water access and financial constraints are the primary impediments to cultivation.

This paper is closely related to the literature on irrigation development, land cultivation, and response to water availability. Duflo and Pande (2007) study the upstream and downstream impacts of dam construction in India, finding that while agricultural productivity increases and sensitivity to weather shocks is reduced downstream, poverty incidence actually rises upstream. Likewise, Strobl and Strobl (2011) demonstrate that while large dams in Africa benefit regions downstream, the adjacent cropland does not experience such gains. Hornbeck and Keskin (2014) examine agricultural production over the Ogallala

¹ These subsections, locally referred to as *Unités de Mise en Valeur*, do not necessarily comprise contiguous plots of land.

Aquifer and nearby counties in the United States, showing that counties with a greater share of land over the aquifer expanded their water-intensive crop areas over time—ultimately facing similar weather-related risks as those specializing in more drought-tolerant crops. More recently, [Blakeslee et al. \(2023\)](#) show that irrigation boosts long-term economic activity in irrigated villages in India, although improvements in agricultural productivity vary widely, while [Tattheago and D'Odorico \(2022\)](#) find that dam construction in Africa tends to favor large-scale agribusiness over small-scale farming. Additionally, [Rufin et al. \(2018\)](#) report that cropping frequencies are higher in dam-irrigated areas compared to rainfed regions across the globe, although inefficient water transfer can limit the potential for intensive cropping. In a similar vein, [Blakeslee et al. \(2020\)](#) document that exogenous well failure in India prompts households to exit agricultural production.

Within this literature, the most closely related paper to ours is [BenYishay et al. \(2024\)](#) who study small scale irrigation program impacts on agricultural and nutritional outcomes in Mali. We utilize a similar event study design, using raw NDVI as an outcome in robustness checks, and obtain a 60% larger effect than their study. One major difference between contexts is that in Mali irrigation extends and supplements water for agriculture in the rainy season, whereas in the Senegal River Valley there is a large effect of irrigation in the non-rainy season. In addition, our focus in this paper is on detecting land usage patterns and correlates of irrigation project success using high frequency observations afforded by satellite data, whereas [BenYishay et al. \(2024\)](#) are primarily focused on household-level outcomes.

Another closely related paper is that of [Asher et al. \(2022\)](#) who study the long-run impacts of irrigation canals in India. As in our study, they combine information on the timing of irrigation infrastructure completion with a highly disaggregated dataset of outcomes, including cultivation rates. However, their analysis of agricultural outcomes focuses on a period long after irrigation canal completion, 30 years at the median. In contrast, we are able to dynamically trace out the effects of irrigation infrastructure over time and by year of completion. Put differently, [Asher et al. \(2022\)](#) study equilibrium outcomes, whereas we can study transition paths. Another notable distinction is the disaggregation level. Because we have cultivation measurements at the 30 m × 30 m pixel-level, we can study within-project dynamics, such as rotations in crop cultivation.

In much of the existing literature, treated areas are defined based on hydrological or topographical criteria, such as being downstream of a dam, adjacent to a canal, or above an aquifer. These spatially defined command areas encompass heterogeneous land uses including rainfed agriculture, human settlements, road infrastructure or terrain unsuited for cultivation due to slope and soil constraints. As a result, irrigation development can only affect a subset of land that is both suitable for cultivation and not yet irrigated. In contrast, we restrict attention to explicitly demarcated irrigation perimeters that are part of planned irrigation development projects, as in [BenYishay et al. \(2024\)](#). This distinction can help explain why we observe substantially larger impact of irrigation infrastructure on cultivation rates relative to extant studies.

Our use of satellite imagery builds on a large literature on remote sensing that seeks to map agricultural outcomes. Researchers have used a variety of vegetation indices as remote-sensing proxies for plant biomass since the late 1960s, with NDVI being the most widely used ([Xue et al., 2017](#)). More recently, satellite imagery has been used to better understand agriculture in low-resource settings where ground-truthed measurements may be sparse and of lower quality ([Burke et al., 2021](#)). [Lobell et al. \(2019, 2020\)](#) show that peak NDVI, the measure we use in this study to detect cultivation, is more strongly correlated with yields measured by full crop cuts than self-reported yields in the case of sorghum in Mali and maize in eastern Uganda. Because NDVI is consistently measured over time by sensors aboard Landsat satellites, we are able to map cultivation at a very high level of resolution for nearly the past four decades.

The approach we take to mapping cultivation is similar to recent work by [Higginbottom et al. \(2023\)](#), who also study expansion of irrigation in the Senegal River Valley using Landsat imagery. They train a random forest model to classify land as irrigated or not, using a dataset of images hand-labeled to outline areas that appear cultivated and find a large expansion of irrigated area, particularly since 2008. Our approach differs from theirs in two ways. First, we use administrative data on the location of irrigation projects, making the cultivation classification task much simpler as we do not need to discern agricultural land from other persistent vegetation such as stands of trees. Second, we apply our dataset to causally estimating the effect of irrigation project completion on land use, a step beyond the work of [Higginbottom et al. \(2023\)](#) made possible by our access to administrative data on irrigation project completion.

More broadly, there is a growing body of economic research which uses remote sensing techniques to measure agricultural outcomes over large areas and long time scales in causal research designs. [Wuepper et al. \(2023\)](#) employ regression discontinuity at international borders around the world to study the effect of economic freedom, as measured by the Fraser Institute, on agricultural productivity using a global dataset of pixel-level observations of annual maximum enhanced vegetation index (EVI), a modified version of NDVI. The aforementioned study by [Asher et al. \(2022\)](#) also uses EVI as a proxy for agricultural productivity which is observable for every settlement in India. [Murillo-Sandoval et al. \(2021\)](#) develop a 31-year panel of land use based on Landsat imagery and using difference-in-differences find that conflict events in Colombia cause the conversion of forest to agriculture. Our study contributes to this literature by demonstrating how large collections of remotely sensed data can be combined with primary data collection to better support rigorous statistical analysis.

The rest of the paper is organized as follows. Section 2 presents the context for this paper. Section 3 presents our data. Section 4 introduces the econometric specification. In Section 5, we present results on land use. Section 6 discusses constraints to land use. Concluding remarks are given in Section 7.

2. Context of study

The Senegal River Valley is an agroecological region located in the northern and eastern parts of Senegal, with the river itself serving as the boundary with Mauritania.² The region plays a crucial role in domestic food production by supplying major crops such as rice and onion, and is broadly specialized in cereals and horticulture.³ Our research is conducted in the delegations of Dagana and Podor, which are located in the Saint Louis region. Online Appendix Table A1 provides a description of crops cultivated by acreage and season.

Due to minimal rainfall, traditional agriculture in the valley has relied on the use of river water, particularly through cultivation in flooded areas after water recession, which benefits from rich alluvial soils and residual moisture. Early projects to control water flow and transition towards an irrigation model began as early as 1925, yet only 6000 ha were irrigated by the time of independence in 1960.

A major shift in water management occurred in the 1980s with the construction of two significant dams. The Diama dam, situated near the river estuary in the lower valley, was built to prevent saltwater intrusion during the dry season, while the Manantali dam in Mali, located in the upper valley, was designed to store water from the Guinean affluents for year-round management. These dams introduced controlled flooding that enabled traditional agriculture to coexist with a

² See [Fig. 1](#) for a map of Senegal and the study region.

³ The Senegal River Valley accounted for 29% of domestic rice production and 31% of domestic onion production in 2022 and 2018 respectively (Direction de l'Analyse, de la Prévision et des Statistiques Agricoles, 2022; Agence Nationale de la Statistique et de la Démographie, 2018).




Fig. 1. Senegal River Valley study areas and irrigated land expansion dates.

Notes: The map shows development units included in the analysis colored by their year of completion, with the histogram along the bottom panel displaying the growth in area equipped for irrigation by year. Water bodies are shown in blue and the three green dots show the location of the capitals of the three departments (level 2 administrative units) in the Senegal River Valley: Saint-Louis, Dagana and Podor. For the purposes of the analysis, development units contained within the department of Saint-Louis are included in the Dagana delegation. The small inset at the top right shows all the departments of Senegal, with the grey rectangle showing the study region. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

progressively developed irrigation system, now covering approximately one third of the agricultural land under traditional flood agriculture and two thirds via irrigation canals.

The Société d'Aménagement et d'Exploitation des Terres du Delta du Fleuve Sénégal (SAED), henceforth the land development authority, has been the primary government body responsible for this transformation, overseeing the construction, operation, and maintenance of irrigation infrastructure across the Senegal River Valley. The irrigation projects analyzed in this study represent predominantly “hard” investments focused on water infrastructure, with minimal integration of complementary institutional or organizational support programs. Initially, large tracts of irrigated land ranging from 500 to 2000 hectares were allocated to small-scale farmers organized at the village level, leading to an increase from 6000 to 18,000 ha under village-based management (Bruckmann, 2018; Le Roy, 2000).⁴ In the early 1990s, the emergence of private, intermediate-sized developments expanded the irrigated area to 65,000 ha by 1994 (Faye et al., 1995), and since 2000, larger private developments and the rehabilitation of older projects—supported by international organizations such as Germany's KfW, the World Bank, Japan's JICA, the United States' MCC, the French Afd, and the African Development Bank (Bourgoin and Diop, 2023)—have raised the total irrigated area to 128,000 ha, out of an estimated irrigable potential of 240,000 ha (SAED and JICA, 2019; MAER, 2015; USAID, 2017). Satellite measurements of cultivated land, shown in Appendix Figure A1, broadly corroborate these official estimates.

Throughout this study, we use two terms to refer to units of agricultural land: “project” and “development unit”. A project covers all plots

that share the same hydraulic infrastructures (e.g. adductors, irrigation perimeters, or drains) or the same water source (e.g. a pumping station, or a water intake on a large canal). In our analysis, a project is classified as public if its management is overseen by the land development authority. Conversely, if the management is handled by independent or private entities, even in joint ventures, the project is categorized as private.

Our basic unit of analysis is a development unit, which corresponds to the land operated by a single producer organization or independent private actors within a project. In our sample, development units contain 44 members on average and projects contain an average of 1.6 development units. The division of projects among development units simply corresponds to an organizational and property rights division, and does not translate into a hydraulic division.

3. Data

Our analysis relies on three datasets. The first is administrative data from the land development authority on development units in the delegations of Dagana and Podor. These data inform the geographic boundaries of development units as well as the characteristics of associated projects and producer organizations. The second is Landsat satellite imagery spanning over three decades that we use to obtain granular high frequency measures of land cultivation by season. The third is a phone survey of farmers in irrigated development units in Podor that project leaders determined to have low or intermittent cultivation. We describe each of these three data sources in what follows.

3.1. SAED administrative data

The sample for our analysis is built using a collection of datasets provided by the land development authority detailing irrigation projects

⁴ In these small-scale public schemes the land development authority assumed a coordinating role by organizing beneficiary farmers into village-level producer groups.

as of 2020. The first dataset of this collection is a shapefile delimiting the boundaries of development units and the associated project. The polygons were mapped in the field by the authority's personnel following the completion of irrigation infrastructure construction or rehabilitation. The project and development unit identifiers contained in the shapefile are used to merge in year of completion of each irrigation project, whether it is publicly or privately managed and the number of members, among other variables.

We limit the analysis sample to projects constructed in 1988 at the earliest to allow for assessment of pre-trends, given that our satellite imagery data begin in 1985. Development units are excluded from the analysis sample if any of the following are true: (i) they cannot be georeferenced via the shapefile (1513 development units), (ii) their construction date is not available in the project characteristics dataset (259 development units), or (iii) they are associated with a project that was constructed before 1988 (796 development units). In total, our sample consists of 1364 development units from 860 different irrigation projects. The development units in our sample cover 46,055 hectares, about 51% of the irrigated area recorded in the shapefiles. The location of the development units in space and their associated project construction dates are shown in Fig. 1.

In order to better understand whether the analysis sample is representative of the broader development units in the study area, Appendix Table A2 compares the means of observable characteristics between development units that are included and excluded from the analysis. The table reveals that the development units in our sample are somewhat more likely to be private and are relatively newer compared to excluded development units. The latter is partly mechanical, due to our removal of projects completed before 1988. In terms of number of members and share of female members, both samples appear similar. Using data from 2021, the development units in our sample have a similar average size compared to excluded development units. Importantly, while Table A2 shows that development units in our analysis sample differ along a number of observable dimensions from those that are excluded, we also show that our analysis sample is quite similar to the set of all irrigated land in the Senegal River Valley in terms of outcomes. Specifically, we construct the 1985–2019 NDVI series for development units in our analysis and for a set of hand-drawn irrigation project polygons compiled by Zwart (2017). Appendix Figure A2 shows that, while there is a slight divergence in the cumulative density functions of NDVI by sample, they are extremely close over much of the range of observed NDVI values, suggesting that our sample is broadly representative in terms of cultivation levels.

3.2. Satellite imagery

We construct cultivation measures using satellite imagery spanning the 1985–2019 period. In particular, we use Tier 1 Thematic Mapper, Enhanced Thematic Mapper+, and Operational Land Imager imagery from Landsats 5, 7, and 8, respectively. These sensors record multispectral imagery data for 30 m × 30 m pixels roughly every two weeks. Landsat 5 was operational from March 1984 to May 2012, Landsat 7 from January 1999 to April 2022, and Landsat 8 launched in April 2013 remains operational.

For the event study analysis, we restrict the temporal coverage of the study to 1985–2015, as the imputation event study estimator we use requires a set of projects that have not yet been completed for identification and all observed projects are completed by 2020. Furthermore, there are very few projects completed after 2016 in our sample, meaning that there are too few controls to estimate treatment effects with meaningful precision in years 2016 or later. There is relatively worse temporal coverage prior to the year 2000, as only Landsat 5 was operational and produced relatively few science-quality

images.⁵ We apply a cloud mask based on the Landsat CFmask and apply the relevant scaling factors to surface reflectance values before any additional processing.

Our analysis is based on a measure derived from multispectral imagery, the Normalized Difference Vegetation Index (NDVI), which has been widely used to study agricultural production due to its strong correlation with biomass. NDVI exploits the fact that chlorophyll reflects electromagnetic radiation with wavelength greater than approximately 0.69 μm , but absorbs visible light (Myneni et al., 1995). This causes plants to have high reflectance values in near-infrared bands, but low values in the adjacent red band.⁶

NDVI is calculated as

$$NDVI = \frac{\text{Near Infrared} - \text{Red}}{\text{Near Infrared} + \text{Red}}$$

with greater values indicative of greater biomass. While the potential range is from -1 to 1 , values less than 0 are rarely observed in our sample.

We calculate NDVI in every available image of the study region between 1985 and 2019 and take the pixel-level maximum, either for each of the three agricultural seasons in a given year (rainy, cold off-season, hot off-season) or combining the two off-seasons (cold off-season and hot off-season). To measure agricultural cultivation, we consider December–March as the cold off-season, April–July as the hot off-season, and August–November as the rainy season.

To get a better sense of how NDVI captures changes in agricultural production, we collect the red, green, and blue (RGB) bands of images over an area with varying NDVI for the 2020 hot off-season, and create a composite image with RGB values of a given pixel associated with the date of maximum NDVI for that pixel. In a sense, this composite shows the area as it would look if every pixel attained its maximum NDVI at the same time. We implement the same procedure in 1984, before irrigation. The left hand side image in Fig. 2 shows the area in 1984. As can be seen, the image shows a barren landscape with no vegetation. Accordingly, maximum NDVI in the perimeters outlined in pink and green is only around 0.1 . In contrast, the right hand side of Fig. 2 shows the same area in 2020. The development unit outlined in bright green has an average NDVI of 0.7 in 2020 and looks dark green in the RGB image, consistent with successful cultivation. In contrast, the development unit outlined in pink, which has the lowest average NDVI in the image at 0.21 , looks brown, consistent with null or unsuccessful cultivation. This ‘smell test’ suggests our use of maximum NDVI works at capturing cultivation patterns from satellite images, as intended.

We next deal with the definition of a cutoff for what we classify as a cultivated plot. We propose a measure for whether a pixel is cultivated based on a simple threshold rule. If a pixel achieves a maximum NDVI in a given season greater than 0.3 , we classify that pixel as cultivated. Our use of the max NDVI value is adequate given that development of biomass over a season may differ between locations due to management decisions, weather, and crop choice. Furthermore, given the sparsity of imagery in the early years of our sample, an ideal classification method would work well with imagery captured at any point during the season. Using a relatively low threshold that is greater than the maximum level generated by background vegetation allows for discrimination of cultivation prior to crop maturity.

The threshold of 0.3 is selected based on the fact that there are very few (less than 5%) pixels in uncompleted projects that exceed this value

⁵ To evaluate the potential influence of temporal variation in image availability on our estimates, we re-estimated our main results using a restricted sample that includes no more than four randomly selected images per year. This restriction is intended to mitigate the disproportionate impact of limited image availability prior to 2000 on estimated cultivation rates. The results remain robust under this alternative specification.

⁶ The near infrared and red bands correspond to $0.77\text{--}0.9 \mu\text{m}$ and $0.63\text{--}0.69 \mu\text{m}$ respectively for the Landsat 5 Thematic Mapper.




Fig. 2. Visualizing NDVI.

Notes: This figure displays an NDVI-priority mosaic of an area in Dagana during the hot off-season in 1984 and 2020. Red, green, and blue pixel values were selected from the image with the greatest NDVI value at that pixel. The development unit with the pink outline corresponds to the lowest NDVI in 2020 in the figure and the bright green to the greatest. The maximum hot off-season NDVI values for these example development units are given by the text in their respective color. Other development units included in the analysis sample are outlined in black. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

during the off-season. Appendix Figure A3 shows that, for the majority of years we study, the 95th percentile of NDVI values in projects which have not been completed is less than 0.3, and it is less than 0.4 in all but one year, suggesting that NDVI rarely reaches above these levels in the absence of irrigation. We also conduct a series of robustness checks using alternative thresholds and raw NDVI demonstrating that results are robust to threshold choice. Development unit-level cultivation rates are calculated by taking the simple average of this binary pixel-level use status.

Figure A4 in the Online Appendix displays the three basic steps to construct our main outcome measure. The top image in the figure displays a sample composite image with true color (RGB) satellite imagery in which the values for each pixel are obtained from the image with the largest NDVI for that pixel. A few things stand out. First, the image shows that background vegetation in this context is virtually non-existent. Multiple field visits confirmed this interpretation — yellow areas are indeed barren and cultivated areas look green in the imagery.

Middle image (b) displays the raw NDVI measure constructed from the image above. Darker green corresponds to higher NDVI. Contrasting the top to the middle image provides the intuition that NDVI indeed captures what appear as cultivated areas well. Finally, bottom image (c) displays our main outcome measure, which is 1 (in dark green) if NDVI in a pixel is above the cutoff of 0.3. Comparing the bottom image to the middle and top images (especially within the black contours which are the development units used in our analysis) suggests our main measure captures land cultivation well.

3.3. Phone survey

To elucidate the reasons for land non-use, we conducted a phone survey with farmers in the delegation of Podor in September 2023. We started our sampling frame with the satellite-derived land use presented in Section 3.2. First, we pre-selected 30 development units in the delegation of Podor that exhibited multiple years of non-use,

based on the estimated land use methodology presented in Section 3.2.⁷ Next, we worked with the land development authority to obtain contact information for representatives of 20 of those development units.

For each of the twenty selected projects, we conducted in-person interviews with a designated representative of the corresponding development unit using a standardized survey instrument.⁸ At the end of these general in-person surveys, we asked each respondent to give us the contact information of 10–15 farmers in their development unit *who have land they do not use, either continuously or intermittently*. In the end, we received the contacts of 206 producers and successfully conducted interviews with 162 of them, obtaining information on 268 different plots.

The phone survey asks about the history of land use in plots operated by the farmer during the three off-seasons of 2021–2023. It also asks about reasons for lack of use, separately for land that is never used in the three off-seasons (continuous non-use) and for land used some of the time (intermittent non-use). Potential responses include water availability, financial constraints, labor availability, and land quality, among others. We also asked respondents whether the main issue constraining use is a problem that is relevant only for the farmer, or whether the problem is relevant to the whole development unit, or the project. We followed up with a question on whether the issue raised could be potentially solved by the farmer himself, by the development unit, or whether it would require external intervention beyond the producer organization.

We ended the survey with a series of preference elicitation questions about constraint-alleviating services that the respondent would be interested in receiving. In line with our constraints of focus, we

⁷ In making this pre-selection, we restricted the sample to development units that are within 5 km from the main highway in order to minimize travel time to the chosen sites.

⁸ This in-person survey is different in content from the phone survey and was mainly fielded to verify the validity of the satellite image land use methodology discussed in Section 3.2.

considered the following six service options: (i) put you in touch with a financial institution, (ii) connect you to someone who is interested in renting out land, (iii) connect you to someone who is interested in renting in land, (iv) provide you with information on how to improve land quality, (v) put you in touch with someone that can provide reliable labor for your farm, or (vi) coordinate a meeting with your neighbors to facilitate improvements to the drainage and irrigation system. We presented farmers with a random selection of 3 of the 15 possible pairwise combinations of these improvement options and ask them to choose which one of the two they would prefer. These light-touch options are all costless to the farmer and do not imply any subsidy. The intent was to glean information about which type of problems are perceived by the farmers as solvable.

4. Econometric specification

We organize our data as a panel of development units, each of which has a potentially different irrigation treatment date (i.e. staggered roll-out). Given this data structure, we implement an event study specification that ensures our estimates have a plausibly causal interpretation. Specifically, we implement the imputation estimator developed by Borusyak et al. (2024). To do this, we first posit a model of potential outcomes in the absence of irrigation:

$$y_{u,t}(0) = \alpha_u + \delta_t + \varepsilon_{u,t} \quad (1)$$

where u indexes development units and t indexes years. This model can be augmented to allow for estimation of pre-trends. Letting $E_{p(u)}$ denote the year the project associated with development unit u is constructed and $pre = \{-9, -8, \dots, -1\} \setminus \{-3\}$, this specification can be written as:

$$y_{u,t}(0) = \alpha_u + \delta_t + \sum_{k \in pre} \beta_k 1\{t - E_{p(u)} = k\} + \beta_{-10} 1\{t - E_{p(u)} \leq -10\} + \varepsilon_{u,t} \quad (2)$$

The indicator for three years prior to project construction is omitted because it is the first that all observations share.⁹ Both Eqs. (1) and (2) are estimated using only observations of development units in years prior to project construction. Once estimated, the coefficients and associated standard errors (clustered at the project level) of Eq. (2) allow for an evaluation of pre-trends to partially assess the validity of the parallel trends assumption underlying causal identification. The model of untreated potential outcomes implies that $\beta_k = 0$ for all values of $k \in (\{-10\} \cup pre)$, which can be formally evaluated via an F-test.

Next, we use the estimated parameters of Eq. (1) to impute untreated potential outcomes for post-treatment observations, $\hat{y}_{u,t}(0) = \hat{\alpha}_u + \hat{\delta}_t$. Post-treatment observation-specific treatment effects can then be simply estimated as $\hat{\tau}_{u,t} = y_{u,t}(1) - \hat{y}_{u,t}(0)$. While these development unit specific estimates are potentially interesting in their own right, performing statistical inference on them is impossible without some restriction of treatment effect heterogeneity and the estimator underlying their computation does not necessarily exhibit desirable properties, such as consistency.¹⁰ Given a set of regularity conditions described by Borusyak et al. (2024), however, averages of the observation-specific treatment effects can be consistently estimated. For this reason, our estimands of interest are averages of observation-specific treatment effects, which can be obtained by assigning weights $w_{u,t}$ to each post-treatment observation and taking the weighted average

$$\hat{\tau}(\{w_{u,t}\}) = \frac{\sum_{u,t} w_{u,t} \hat{\tau}_{u,t}}{\sum_{u,t} w_{u,t}} \quad (3)$$

⁹ Under the proposed model of untreated potential outcomes, the choice of which event time indicator(s) to omit is asymptotically irrelevant, as all estimated pre-treatment coefficients will converge to 0. Omitting an event time indicator earlier than -1 , the event time which is typically omitted, allows for more straightforward visual inspection of pre-trends in the periods immediately prior to treatment, when anticipation effects are most probable.

¹⁰ This is because in short panel settings, the unit fixed effects α_u are not consistently estimated.

where $w_{u,t} = 0$ for pre-treatment observations. To estimate typical event study effects, we set $w_{u,t} = 1\{t - E_{p(u)} = k\}$ for various values of $k \geq 0$. Borusyak et al. (2024) also provide a method for conservative estimation of clustered standard errors of $\hat{\tau}(\{w_{u,t}\})$. We extend their statistical inference procedure to allow for clustering on irrigation projects, rather than development units, given that this is the level at which treatment is assigned.

Aside from avoiding potential biases due to negative weighting that can arise when using the more traditional two-way fixed effects estimator in empirical settings such as ours with staggered treatment timing and (potentially) heterogeneous treatment effects, the imputation estimator confers advantages over other recently-proposed estimators. In particular, our panel is not complete, with observations much more likely to be missing earlier in the panel, rather than later. These missing data pose no identification issue for the imputation estimator, so long as we are able to estimate $\hat{\alpha}_u$ and $\hat{\delta}_t$. This is because once $\hat{\alpha}_u$ and $\hat{\delta}_t$ are estimated, we can calculate $\hat{\tau}_{u,t}$ for any development unit-year for which we observe the treated outcome. Under the proposed model of untreated potential outcomes, additional pre-treatment observations only improve the precision with which $\hat{\alpha}_u$ and $\hat{\delta}_t$ are estimated, and thus the precision of $\hat{\tau}(\{w_{u,t}\})$.

5. Impact of irrigation project completion on land use

5.1. Aggregate effect

We begin by presenting preliminary evidence of the impact of irrigation project completion in the raw data. Fig. 3 displays monthly averages of NDVI in 2015, a year for which we have an abundance of images that allow this type of month-level analysis, for development units in projects completed in 2014 or earlier compared to those completed in 2016 or later. In the figure, the orange shaded box corresponds to the hot off-season, the purple shaded box to the rainy season, and the white area to the cold off-season. The dashed orange line is average NDVI in completed projects, while the solid green line displays average NDVI among projects that have yet to be completed.

Three facts emerge from this raw data figure. First, it is striking that measured NDVI is uniformly higher in completed projects throughout the year, including in the rainy season. Second, the figure also shows that the gap in NDVI is largest in the hot off-season, when NDVI is almost double in complete vs. incomplete projects. Lastly, the pattern over the course of the year suggests that, for pre-completion projects, agricultural production is concentrated in the rainy season, whereas for post completion observations NDVI is at high levels in both the hot off-season and the rainy season, consistent with irrigation allowing for two agricultural cycles in a year.

Our main results for the impact of irrigation project completion on land cultivation rates are presented in Fig. 4. The figure displays coefficient estimates and 95% confidence intervals of the event study estimation procedure described in Section 4 with the NDVI threshold-based measure of development unit cultivation rate as the outcome. Recall that the outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI over the season (either August–November for the rainy season, or December–July for the off-season) is greater than 0.30, and 0 otherwise. The values of these pixel-level indicators are then averaged at the development unit level to obtain the share of project land in use. The X-axis shows the number of years since the project creation date, which is the first year of operation as recorded by the land development authority. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level.

First, the figure provides compelling evidence that there are no pre-trends in cultivation rates based on the pattern of coefficients that are tightly centered around zero up to 10 years before the project begins.




Fig. 3. NDVI by month before and after project completion.

Notes: This figure displays the monthly average maximum NDVI series by construction status of the project for the year 2015. Development units in the “Pre” group were completed in 2016 or later whereas development units in the “Post” group were completed in 2014 or earlier. The orange shaded box corresponds to the hot off-season. The purple shaded box corresponds to the rainy season. The white area for December–March corresponds to the cold off-season. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 4. Effect of irrigation project completion on cultivation rate.

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the event study estimation procedure with the NDVI threshold-based measure of cultivation rate as the outcome. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the development unit level over the period August–November (Rainy Season), and the rest of the year (Off-season) to get a proxy for the share of development unit land in use. Average cultivation rates prior to project completion are 0.097 (Rainy season) and 0.0402 (Off-season) across the sample. The x -axis shows the number of years since the project completion date, which is the first year of operation as recorded by the land development authority. The red vertical line separates the treatment leads from the treatment lags. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from standard errors clustered at the project level for pre-trend coefficients and conservative estimates of standard errors clustered at the project level for post-treatment coefficients. The sample is restricted to the 1364 development units in projects constructed in 1988 or later which could be merged with the shapefile. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 5. Effects of irrigation project completion by year.

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the effect of irrigation project completion on off-season cultivation rate by year. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the development unit level over the period December–July to get a proxy for the share of development unit land in use during the off-season. The average cultivation rate prior to project completion is 0.0402 across the sample. For each year s , we set the weights $w_{u,t} = 1\{t - E_{p(u)} > 0\} \times 1\{t = s\}$ for values of $s \in \{1988, \dots, 2015\}$, generating estimates of the average effect of irrigation in s . The x-axis shows the year for which the effect is estimated. The black dots are coefficient estimates while the black vertical lines are 95% confidence intervals constructed from conservative estimates of standard errors clustered at the project level. The sample is restricted to the 1364 development units in projects constructed in 1988 or later which could be merged with the shapefile.

Second, cultivation rates start picking up substantially beginning in the first year after irrigation completion. Cultivation rates are remarkably stable at around 25 percentage points higher in the treated development units for the first twenty years. This represents a seven-fold increase from 4 to 29 percentage points in the non-rainy season and a tripling from 10 to 35 percentage points cultivated in the rainy season.

Beyond twenty years after irrigation project initiation, the estimated effects continue to rise, particularly during the off-season, with cultivation rates exceeding pre-irrigation levels by more than 50 percentage points. These are very large effects, considering that average cultivation rates prior to project completion were only 10 percent in the rainy season and 4 percent in the off-season across the sample, implying an increase of about an order of magnitude in land cultivation rates twenty-five years post project completion in the off-season. In the rainy season, cultivation rates are more than four times greater than at baseline by the end of the analysis period. Overall, these findings refute the hypothesis that irrigation interventions yield merely transient gains that subsequently dissipate. Instead, impacts on cultivation rates are large and sustained.¹¹

We also perform robustness checks with alternative thresholds of 0.4 and 0.5 as well as raw NDVI. The results, shown in Appendix Figures A6–A8, are qualitatively similar, though effects are mechanically smaller in magnitude for greater thresholds. Our estimates for raw NDVI are about 60% larger than those in BenYishay et al. (2024), which are directly comparable to ours. One difference in contexts is

that irrigation in Mali is used to extend and supplement agriculture during the rainy season, whereas in Senegal it allows for cultivation in the off-season.

We can also consider average impacts over time for all completed projects. To do so, we set $w_{u,t} = 1\{t - E_{p(u)} > 0\} \times 1\{t = s\}$ for values of $s \in \{1988, \dots, 2015\}$, generating estimates of the average effect of irrigation in year s .

Fig. 5 presents the estimated effects of irrigation on off-season cultivation, disaggregated by calendar year. In contrast to the event study results – which show relatively stable effects by years-since-completion – we observe substantial heterogeneity in irrigation impacts across calendar years. Specifically, the effects on off-season cultivation are modest prior to 2000 but increase steadily thereafter. The impact reaches its peak in 2015, with an estimated effect just below 50 percentage points. Changes in the availability of satellite imagery – which occurred in our sample in 1999, 2012, and 2013 – as well as algorithm processing changes can lead to measurement error issues (Jain, 2020; Wuepper et al., 2025). We test for these issues by including development unit $\times$ satellite fixed effects (Appendix Figure A9) and by restricting our imagery to up to four randomly selected images in a given development unit per season (Appendix Figure A10), and find qualitatively similar and less-precisely estimated results.

To examine the role of irrigation in facilitating climate adaptation, we estimate the effects of seasonal averages of daily minimum and maximum temperatures on NDVI, both before and after the completion of irrigation projects. We find that prior to irrigation, minimum and maximum temperatures have significant and opposing effects on NDVI, consistent with previous research on temperature sensitivity in rice production (Welch et al. 2010). However, these effects largely dissipate after irrigation is introduced (Fig. 6). This pattern suggests that, in addition to increasing overall cultivation, irrigation substantially reduces vulnerability to temperature shocks—an important channel

¹¹ For this analysis, we considered as a single season the hot and the cold off-seasons. We do this because farmers engage in at most one cultivation cycle in the off-season, either cold or hot but not both, due to temporal overlap between harvesting and land preparation activities. Appendix Figure A5 presents results for each off-season separately, confirming that the effects remain evident in both the cold and hot off-seasons.




Fig. 6. Effects of temperature by presence of irrigation.

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the effects of average daily minimum and daily maximum temperature on NDVI by irrigation project completion status and agricultural season. Each of the four panels contains estimates from a separate regression. Temperature data comes from the ERA5 Reanalysis. Additional controls include total precipitation over the season, development unit fixed effects and delegation-year fixed effects. The specification is given by $NDVI_{u,t} = \alpha_u + \delta_t + \beta_{min} Min Temp_{u,t} + \beta_{max} Max Temp_{u,t} + \gamma Precip_{u,t} + \varepsilon_{u,t}$. The sample is restricted to years 2000 and later to ensure that changes in the level of NDVI are not due to different availability of imagery. Standard errors are clustered at the project level.

through which irrigation contributes to climate resilience. These findings align with recent evidence from the climate and agriculture literature (e.g., Wang et al., 2024).

5.2. Heterogeneity in patterns of use across development units

We next explore whether effects are heterogeneous across cohorts. We explore this question in Fig. 7, separately estimating treatment effects by year for cohorts of projects completed between 1988–1999, 2000–2009, and 2010–2015. While the cohort of projects completed between 2000 and 2009 does perform better than older projects at the turn of the century, the oldest projects still experience a jump in cultivation in 2000 and consistent increase thereafter. These results suggest that, while there is some cohort effect, much of the heterogeneity we observe by year holds across all projects.

Next, we assess whether management type (public or private) predicts differential impacts of irrigation. Estimating treatment effects by year for development units in public and private projects separately shows that public projects generally perform better, with the starkest difference in the period from 2000 to 2007 (Fig. 8). After 2008, there is a substantial convergence, though public projects still outperform private projects in the majority of years.

Finally, our data allow us to examine whether cultivation rates vary systematically with the extent of nearby irrigation development, which could create competition for water resources. To investigate this, we estimate year-five treatment effects at the development unit level and calculate the total area newly developed for irrigation within concentric doughnuts of 1 km to 5 km around each unit. We then regress these year-five treatment effect estimates on the extent of new irrigation development within each doughnut. If surrounding units divert water away from the development unit, we would expect greater nearby irrigation expansion to be associated with smaller treatment effects. However, we find little evidence of such negative spillovers. While the estimated coefficients are negative, their magnitudes are negligible (See Appendix Figure A11).¹²

5.3. Development unit-level land use classification

The results presented in Sections 5.1 and 5.2 show that irrigation yields substantial impacts on average, as well as heterogeneous

¹² We also find no evidence of systematic heterogeneity in treatment effects as a function of a project's location downstream along the Senegal River.




Fig. 7. Effects of irrigation project completion by year and Cohort.

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the effect of irrigation project completion on off-season cultivation rate by year and project cohort. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the development unit level over the period December–July to get a proxy for the share of development unit land in use during the off-season. The average cultivation rate prior to project completion is 0.0402 across the sample. The x-axis shows the year for which the effect is estimated. The dots are coefficient estimates while the vertical lines are 95% confidence intervals constructed from conservative estimates of standard errors clustered at the project level. Colors and shapes represent project cohorts, with green circles indicating projects completed between 1988 and 1999, orange triangles between 2000 and 2009, and blue squares between 2010 and 2015. The sample is restricted to the 1364 development units in projects constructed in 1988 or later which could be merged with the shapefile. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

responses across development units. In this section, we characterize this heterogeneity more generally by classifying individual development unit-level time series into four use classes associated to different degrees of average cultivation rates post project completion.

Fig. 9 shows one randomly selected example of each class that is associated with a project completed between 2000 and 2005. The top left panel is an example of a “High Use” project development unit, defined as one in which average cultivation rates exceed 67 percent in the post-completion period. In this particular example, cultivation rates shoot up dramatically immediately after the project is finalized – to around 87 percent – and in some years reaching close to 100 percent cultivation rates.

We also observe what we call “Medium Use” cases, in which cultivation rates are between 33 and 67 percent on average post project completion. In the top right panel of the figure, we observe an example that exhibits a more gradual increase in the cultivation rate, reaching 70 percent by year 3, but then stabilizing for a few years. Beginning in 2010, there is a large and persistent fall in cultivation, thus dragging down the average.

The bottom left panel shows an example of a “Low Use” case, defined as those development units with an average cultivation rate lower than 33 percent. In this example, the low average post completion cultivation rate is due to a very slow ramp-up and a low ceiling. More than 10 years post completion, cultivation rates are below 30 percent and they almost never pass the 50 percent mark in over 20 years after irrigation is introduced.

Finally, the bottom right panel displays an example of “Intermittent Use”. Here, we observe the very high cultivation rates post completion, but also years with low or even zero cultivation rates in an alternating pattern. To generalize this concept quantitatively, for each

development unit we calculate the average share of pixels that change cultivation status (from zero to one and vice versa) between years $t - 1$ and t in post-project completion years. We use this metric to classify projects as having intermittent land use if cultivation rates vary by more than 10 percent annually on average, excluding years where either observation is missing.

In order to show the overall prevalence of these different use classes, Fig. 10 plots development units according to their average share of land used and the share of pixels which change cultivation status in the post project completion off-seasons. Here we maintain the definition of “Intermittent Use” with at least 10% of pixels changing off-season cultivation status in the average year, and classify the other development units into “Low Use” as those development units with an average share in use below 33 percent, “Medium Use” with between 33 and 67 percent, and “High Use” with above 67 percent. In the figure, each circle size is proportional to the development unit’s land area. We observe that 18 percent of land can be classified as Low Use, 37 percent as Intermittent Use, 29 percent as Medium Use, and only 16 percent as High Use.¹³ Measured in terms of either units or land area, the intermittent use class is the most prevalent land use category. High use development units represent a smaller share of land than of number of units, showing that successful high-use development units tend to be relatively small in total size.¹⁴

¹³ In Appendix Figure A12, we show the spatial dispersion of all four use class categories. The figure shows that all four different use classes are scattered throughout the analysis area and not concentrated in particular sections.




Fig. 8. Effects of irrigation project completion by year and sector.

Notes: This figure displays the coefficient estimates and 95% confidence intervals of the effect of irrigation project completion on off-season cultivation rate by year and whether the project was publicly or privately developed. The outcome is constructed from a pixel-level binary indicator, which takes the value of 1 if the highest value of NDVI is greater than 0.30 at the pixel level, and 0 otherwise. The values of these pixel-level indicators are then averaged at the development unit level over the period December–July to get a proxy for the share of development unit land in use during the off-season. The average cultivation rate prior to project completion is 0.0402 across the sample. The x-axis shows the year for which the effect is estimated. The dots are coefficient estimates while the vertical lines are 95% confidence intervals constructed from conservative estimates of standard errors clustered at the project level. Colors and shapes represent the sector which developed the projects, with green circles indicating private development, orange triangles public development. The sample is restricted to the 1364 development units in projects constructed in 1988 or later which could be merged with the shapefile. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We can also estimate the main event study figure for the four different land use classes, shown in Appendix Figure A14. Not surprisingly, off-season impacts of project completion are largest for the high use category, and lowest for the low use category. Intermittent and medium use class types track each other closely in terms of impacts on cultivation rates over time.

The finding that 37 percent of all land presents intermittent usage patterns is novel and is a direct result of our satellite-derived high-frequency land cultivation panel at the development unit level. In Online Appendix—Land Rotation, we explore whether intermittent usage is in fact due to land fallowing that allows the land to recover before being used again. We argue that efficient fallowing in a development unit should take the form of rotation of fields in and out of cultivation while keeping overall cultivation shares steady for output and consumption smoothing purposes. We do not find evidence that the intermittent cultivation we observe is due to fallowing.

5.4. Predictors of use class

We documented that development units broadly fall into four classes of land use patterns: low, intermittent, medium, and high use. A

¹⁴ Because cultivation rates have changed so much over time, it is possible that some development units would have a different use class if analyzed using only more recent data. Appendix Figure A13 shows the distribution of use classes when only using data for 2010–2019. While there is an increase in the share of high use development units, this is almost entirely driven by reductions in low and medium use. The share of development units and land that is intermittently used has remained approximately constant, even when only using more recent data.

natural question that arises is whether any characteristics of projects or producer organizations are associated with particular use classes. To explore this question, we focus on the most recent 10 years of data which display high levels of usage and allows a comparison across characteristics that is not confounded with early period observations characterized by low usage rates. Table 1 shows the distribution of development units across the four use classes by three sets of characteristics: cohorts, private vs. publicly managed, and whether the development unit is an agribusiness.

Panel A shows that development units from the oldest project cohort are most likely to be low use, whereas the other cohorts are more likely to be in the other use classes. Notably, all three cohorts are mainly distributed between the intermittent and high use classes, suggesting that general improvement in the effects of irrigation over time has led to generally higher levels of cultivation, either constantly or sporadically.

Section 5.2 showed that publicly developed projects tend to experience greater benefits from irrigation in terms of off-season cultivation rates than private projects. This finding is echoed in Panel B, which shows that development units in private projects have higher shares among the low use type (3 times larger share than public projects) and have 1/3 lower prevalence among the high use class. So even when we focus on the last ten years of data, public projects continue to outperform private ones in terms of land usage.

Panel C shows that agribusinesses are clearly mostly classified as high use (3 out of 4 cases), whereas non-agribusinesses are only classified as high use in about 1 out of 3 cases. We examined whether agribusinesses disproportionately entered areas that exhibited higher cultivation rates among non-agribusiness units prior to their entry, either in levels or in trends. However, we find no evidence of such a




Fig. 9. Development unit-level cultivation series examples.

Notes: This figure displays the off-season cultivation rate time series for four development units. The vertical red line in each panel is the creation date of the project the development unit is associated with. Panels a-d correspond to high, medium, low, and intermittent use classes respectively. For each use class, the example was selected randomly from the set of development units in projects completed between 2000 and 2005. “Low Use” are development units with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Medium Use” are development units with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “High Use” are development units with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Intermittent Use” are development units with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

pattern, which would have indicated selection on the outcome variable. We also investigated whether the entry of agribusinesses adversely affects outcomes in neighboring projects but again found no robust evidence to support this hypothesis.

5.5. Contributions to uncultivated land

While the land use rate has rapidly increased since 2000, unused land remains a major issue, with around 25% of total land in completed projects being unused in 2019 in both the dry and rainy seasons (Appendix Figure A15). We have documented that development units can be classified based on their cultivation rate and share of land that changes cultivation status in a typical year, and that there are

a number of characteristics which predict these use patterns. One question that remains is how much each use class contributes to the total uncultivated area which could feasibly be used for production. We can directly measure this number for development units in our analysis sample by calculating the share of pixels with a maximum off-season NDVI less than 0.3, multiplying by the area of the development unit, and aggregating by use class.

Fig. 11 shows the total off-season uncultivated area of development units within completed projects by use class and year. Notably, in almost every year development units in the intermittent use class contribute the greatest amount of uncultivated land. Medium use development units also contribute about the same amount of uncultivated land each year as low use development units. Between 2000 and 2019, about 11,234 ha per year, or one third of the total land in completed




Fig. 10. Definition of use classes.

Notes: This figure displays development units colored by their use class. The x-axis represents the average share of land cultivated during the off-season in years after project completion. The y-axis shows the average share of land that changes cultivation status during the off-season between the previous and current year for years after project completion. “Low Use” are development units with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Medium Use” are development units with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “High Use” are development units with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Intermittent Use” are development units with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year. Area of circles is proportional to the area of the development unit represented by the circles. The shares reported in the legend are based on area. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

projects, was left uncultivated in intermittent and medium use development units alone. These findings suggest that substantial increases in cultivation and output could be achieved by targeting interventions to increase land utilization in areas with some evidence of use.

6. Causes for non-use or intermittent usage of land

In what follows, we present results from our dedicated farmer survey that allows exploring the bottlenecks holding back over 50 percent of farmland in the Senegal River Valley in the low and intermittent land use categories. Recall that the phone survey sampled farmers that were known in their community for having some land either intermittently or continuously unused. We asked farmers about the history of land use in their plots during the three off-seasons of 2021–2023.

We partition farmers’ plots into three different categories: (i) plots with continuous non-use between 2021 and 2023, which are plots that were not used at all in the three off-seasons, (ii) plots with intermittent non-use between 2021 and 2023, which are plots used in some but not all of the three off-seasons and (iii) plots with continuous use between 2021 and 2023, which are plots used in each of the three off-seasons. In

total, we observe 76 plots continuously unused, 126 plots intermittently unused, and 66 plots continuously used between 2021 and 2023.

Panel A of Table 2 presents aggregate statistics for the constraints to land use analysis. Among our sample of farmers with some unused land, we find 45.6 percent of land intermittently unused (59.2 hectares), 31.5 percent of land continuously unused (40.9 hectares), and 22.9 percent of land continuously used (29.7 hectares). Plots in each category are on average about the same size—half a hectare.

In Panel B of Table 2, we present the main constraints to land use by type of constraint separately for continuously unused and intermittently unused land. This breakdown is useful as continuously unused land (column 1) may well be completely unfit for agriculture. Intermittently used land, on the other hand, is certain to be viable for agriculture and inquiring about reasons for its non-use is exactly what we want to shed light on. The percentages shown in each column represent the area-weighted share of plots that cited the constraint in the left column as their main constraint to land cultivation. These percentages thus represent area shares and they add up to 100 in each column.

Note, first, that in spite of all these plots being in projects with completed irrigation infrastructure, inadequate access to water is still




Fig. 11. Uncultivated area by use class.

Notes: This figure shows the total uncultivated area by use class and year. For each year and development unit, we calculate the cultivation rate and subtract it from unity. This rate is then multiplied by the area of the development unit and summed across development units within a use class $\times$ year to get total uncultivated area for the use class. “Low Use” are development units with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Medium Use” are development units with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “High Use” are development units with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Intermittent Use” are development units with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

Table 1
Shares of development unit by Use Class: 2009–2019.

	Use class				N
	Low	Intermittent	Medium	High	
<i>Panel A: Cohort</i>					
Completed 1985–1999	0.11	0.42	0.13	0.35	800
Completed 2000–2009	0.09	0.39	0.12	0.40	234
Completed 2010–2019	0.05	0.33	0.09	0.53	330
<i>Panel B: Project type</i>					
Private	0.12	0.40	0.13	0.34	795
Public	0.04	0.38	0.10	0.48	569
<i>Panel C: Agribusiness</i>					
Non-Agribusiness	0.10	0.42	0.12	0.36	1203
Agribusiness	0.01	0.19	0.07	0.73	161

Notes: This table displays the share of development units with a given characteristic in each use class. The sum across columns may not necessarily sum to exactly 1 due to rounding. “N” represents the number of development units in the group described in the first column of the table. “Low Use” are development units with an average share in use below 33 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Medium Use” are development units with an average share in use between 33 and 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “High Use” are development units with an average share of use above 67 percent, and less than 10% of pixels changing off-season cultivation status in the average year. “Intermittent Use” are development units with at least 10% of pixels changing off-season cultivation status in the average year, regardless of usage level in the average year.

cited as the main constraint to land cultivation. This is true for 61.6 percent of continuously unused land and 23.9 percent of intermittently unused land.¹⁵

Farmer interviews revealed that irrigation canals are not always able to provide water at sufficient levels for cultivation. This can be due to lack of canal maintenance or to land that is far away from the canal. An example of this is the case of a development unit whose water canals have been ruptured for 5 years, preventing the owners from using their land for cultivation ever since.

In the case of intermittent non-use due to lack of water, one farmer reported that the main electric pump that brought water to the canals of his development unit did not have sufficient capacity to serve all the plots. As a result, the farmers agreed to ration the water so that he would only cultivate his plot in years in which it was his turn to receive water.

Financial constraints were found to be quantitatively more important drivers of intermittent non-use of plots (49.6 percent of area) compared to continuous non-use (25.9 percent of area). In one interview, farmers in a large development unit said they were severely financially constrained, to the point where they could not afford to get a pump, despite close proximity to a major canal. For another

¹⁵ Stars in column (2) of Table 2 represent the significance level associated with the coefficient ρ in the regression of the form: $Y_i = \alpha + \rho X_i + \epsilon_i$, where Y_i is the variable on the first column of the table, X_i is a dummy variable that takes the value of 1 if the plot is intermittently not used and a value of 0 if it is continuously not used, with standard errors clustered at the farmer level. The sample covers only plots of respondents interviewed in the phone survey. Each observation in the analysis corresponds to a plot in the survey.

Table 2
Main constraints to land use.

	Continuous non-use (1)	Intermittent non-use (2)	Continuous use (3)
Panel A: Aggregate numbers			
Number of plots	76	126	66
Share of plots (%)	28.36	47.01	24.63
Total area (ha)	40.90	59.21	29.65
Share of total area (%)	31.52	45.63	22.85
Average area of plot (ha)	0.54	0.47	0.45
Panel B: Average share of constraints weighted by land size (%)			
Waterways not well constructed	23.5	11.9**	–
Low water level at source point	38.1	6.5***	–
Other water constraints	0.0	5.4**	–
<i>Total water constraints</i>	<i>61.6</i>	<i>23.9***</i>	–
No functional pump	8.6	17.0	–
No access to credit	8.3	17.8*	–
High cost of credit	0.0	0.7	–
No desire to borrow	8.3	8.5	–
Other financial constraints	0.7	5.5**	–
<i>Total financial constraints</i>	<i>25.9</i>	<i>49.6***</i>	–
Infertile/unsuitable land	4.2	0.7	–
Soil drainage issues	1.2	4.4	–
Renovation of projects	1.0	12.2**	–
Labor constraints	0.0	2.9**	–
Input constraints	1.5	1.6	–
Other constraints	4.6	4.7	–
<i>Total other constraints</i>	<i>12.5</i>	<i>26.6**</i>	–

Notes: This table shows the main constraints separately for the different type of non-use in the sample of the phone survey. “Continuous Non-Use” in column (1) refers to plots that were entirely not used by respondents between 2021 and 2023. “Intermittent Non-Use” in column (2) refers to plots with land used only for one or two of the three off-seasons between 2021 and 2023. ‘Continuous Use’ refers to plots that were entirely used for each off-season between 2021 and 2023. The shares in each column represents the share of plots that cited the constraint in the first column as their main constraint. The shares are weighted by the size of the constrained areas and they add up to 100 for each column. Stars in column (2) represent the significance level associated with the coefficient ρ in the regression of the form: $Y_i = \alpha + \rho X_i + \epsilon_i$, where Y_i is the variable on the first column of the table, X_i is a dummy variable that takes the value of 1 if the plot is intermittently not used and a value of 0 if it is continuously not used, with standard errors clustered at the farmer level. The sample covers only plots of respondents interviewed in the phone survey. Each observation in the analysis corresponds to a plot in the survey.

* p -value < 0.1.

** p -value < 0.05.

*** p -value < 0.01.

development unit, land was never used due to the lack of a water pump capable of bringing sufficient water to the fields. They mentioned that diesel pumps would be adequate but they did not have the means to afford them.

In theory, functioning credit markets could alleviate some of these financial constraints to deal with issues related to pumps. In practice, however, many farmers reported having trouble accessing financing. Many farmers reported not knowing how to approach getting financing on their own. Others knew how but had trouble finding credit. For example, farmers from one development unit reported having difficulties finding financing for the last few years for a large public project, to the point where they resorted to a wealthy private processor that gave them credit at a heavy premium (in the form of rice bags at harvest). Finally, some farmers mentioned having no desire to borrow for religious reasons or for fear of not being able to pay back. Farmers from one development unit reported having only one pump for their entire land, although the pump does not have the capacity to supply all of the plots. They do not take bank loans to deal with this issue because they are afraid of not being able to repay them.

The ‘Total other constraints’ category was significantly higher for intermittently non-used land. These included subcategories ‘Projects that are being renovated’, as well as problems with ‘access to labor’.

Overall, interviewing farmers with some unused land revealed several new insights. First, fallowing was never cited as a reason for non-use, consistent with result of the rotation analysis presented in the Online Appendix. Second, and more surprisingly to us, water access continues to be a major constraint to land use, especially for continuously unused land. Third, financial constraints were also important causes of non-use, especially among intermittently unused plots.

After asking farmers about the main cause that explained lack of land cultivation in a plot, we asked whether the issue was one that affected the farm/household only, or whether the problem was common to the whole development unit, or was an issue external to the development unit. For example, if lack of water was identified as being the culprit for non-usage of a plot, the farmer was asked to explain if lack of water was due to himself not having brought available water from a nearby canal to his plot (a farm level problem), or whether the development unit was not sending water to the canal serving the farmer (a development unit level problem), or whether it was due to some reason external to the development unit. As an example of the latter, farmers cited the main canal serving the whole development unit not being supplied with enough water. Further, we also asked if the “solution” to the identified constraint was something addressable by the farmer himself, the development unit managers, or would have to be addressed by actors external to the development unit to be solved (for example by the land development authority).

This analysis is presented in Table 3 for continuously unused land. Recall that for these lands, water access was identified as the most important explanation for non-use. What this table reveals is that, for the most part, the water access problem causing non-use is not seen as something the farmer can address on his own. In almost all cases, the solution is seen as needing to be addressed by actors external to the development unit (Column 6, first and second rows).

In Table 4, we present the analogous analysis for intermittently unused plots. Recall that for these types of plots, financial constraints were quantitatively the most frequent reason for non-use. In this case, both the level of the problem and the solution tend to be at the level of the development unit itself.

The second most important cause for non-use among intermittently unused plots is water access. Table 4 (top panel) shows that for plots intermittently unused due to water constraints, the level of the problem and possible solutions are again not considered something that the farmer himself can address. Rather, both the problem and the solution are more likely to be at the development unit level, or external to the development unit. Since this refers mostly to canals not being well constructed from the start or needing major maintenance interventions, we identify water access issues as something the land development authority could potentially improve to increase land usage, addressable via improved canal maintenance and infrastructure projects.

To assess the reliability of farmers’ self-reported responses, we examine the correlation of reported constraints across farmers within the same project. Table A3 reports the variance of reported constraints for the projects with more than one respondent. Overall, the variance is low for both water-related and other constraints (both have a median of 0), indicating that farmers within each project provide largely consistent responses for these issues. In contrast, the variance for financial constraints is higher (with a mean of 0.12 and a median of 0.11), suggesting less agreement among farmers regarding financial issues, or a greater diversity of financial situations within the same projects. In general, the low observed variance within projects aligns with the notion that the responses of farmers are reliable.

6.1. Revealed preference for different constraint-alleviating services

We now present results from our survey module asking questions about potential constraint-alleviating services that the respondent could be interested in receiving at zero cost. We asked farmers to express preference between two of the following options selected at random:

Table 3
Level of problems and solutions for land continuously unused.

	Level of problem			Level of solution			% Area
	Farm/HH	UMV	External	Farm/HH	UMV	External	
	(1)	(2)	(3)	(4)	(5)	(6)	
Water constraints							
Waterways not well constructed	0.0	36.4	63.6	0.0	4.5	95.5	23.5
Low water level at source point	0.00	24.0	76.0	0.0	0.0	100.00	38.1
Other water constraints	–	–	–	–	–	–	0.0
Financial constraints							
No functional pump	0.0	25.0	75.0	0.0	12.5	87.5	8.6
No access to credit	25.0	25.0	50.00	0.0	25.0	75.0	8.3
High cost of credit	–	–	–	–	–	–	0.0
No desire to borrow	100.0	0.0	0.0	20.0	40.0	40.0	8.3
Other financial constraints	0.0	0.0	100.0	0.0	0.0	100.0	0.7
Other constraints							
Infertile/unsuitable land	0.0	25.0	75.0	0.0	0.0	100.0	4.2
Soil drainage issues	0.0	100.0	0.0	0.0	0.0	100.0	1.2
Renovation of projects	0.0	0.0	100.0	0.0	0.0	100.0	1.0
Labor constraints	–	–	–	–	–	–	0.0
Input constraints	0.0	100.0	0.0	0.0	0.0	100.0	1.5
Other constraints	66.7	33.3	0.0	33.3	66.7	0.0	4.6

Notes: This table shows the level of the problem and the level of the potential solution for land continuously not used, which includes all plots that were entirely not used by respondents between 2021 and 2023. The first column shows the main constraint reported by the respondent for the plot-season pair. Columns (1)–(3) show the level at which the constraint binds. Columns (4)–(6) show the level at which the respondent expects the potential constraint alleviation tool to be at. Column (7) shows the share of total area subject to the constraint. Each cell represents the share of total land that falls under the level of problem/solution specified in the corresponding column. Numbers in the same row of columns (1), (2) and (3) add up to 100. Numbers in the same row of columns (4), (5) and (6) add up to 100. The sample is restricted to land continuously not used from Table 2.

Table 4
Level of problems and solutions for land intermittently unused.

	Level of problem			Level of solution			% Area
	Farm/HH	UMV	External	Farm/HH	UMV	External	
	(1)	(2)	(3)	(4)	(5)	(6)	
Water constraints							
Waterways not well constructed	5.6	36.1	58.3	0.0	0.0	100.0	11.9
Low water level at source point	4.2	20.8	75.0	4.2	4.2	91.7	6.5
Other water constraints	0.0	71.4	28.6	0.0	0.0	100.0	5.4
Financial constraints							
No functional pump	0.0	81.8	18.2	0.0	68.2	31.8	17.0
No access to credit	9.4	75.0	15.6	0.0	59.4	40.6	17.8
High cost of credit	0.0	100.0	0.0	50.0	50.0	0.0	0.7
No desire to borrow	57.1	35.7	7.1	14.3	53.6	32.1	8.5
Other financial constraints	27.8	62.2	10.0	10.0	35.0	55.0	5.5
Other constraints							
Infertile/unsuitable land	0.0	0.0	100.0	0.0	0.0	100.0	0.7
Soil drainage issues	0.0	0.0	100.0	0.0	0.0	100.0	4.4
Renovation of projects	5.4	41.1	53.6	0.0	12.5	87.5	12.2
Labor constraints	75.0	18.8	6.2	31.2	62.5	6.2	2.9
Input constraints	25.0	75.0	0.0	16.7	25.0	58.3	1.6
Other constraints	13.6	86.4	0.0	13.6	77.3	9.1	4.7

Notes: This table shows the level of the problem and the level of the potential solution for land intermittently not used, which includes all plots with land used only for one or two of the three off-seasons between 2021 and 2023. The first column shows the main constraint reported by the respondent for the plot-season pair. Columns (1)–(3) show the level at which the constraint binds. Columns (4)–(6) show the level at which the respondent expects the potential constraint alleviation tool to be at. Column (7) shows the share of total area subject to the constraint. Each cell represents the share of total land that falls under the level of problem/solution specified in the corresponding column. Numbers in the same row of columns (1), (2) and (3) add up to 100. Numbers in the same row of columns (4), (5) and (6) add up to 100. The sample is restricted to land intermittently not used from Table 2.

(i) put you in touch with a financial institution (revealing an interest in actions to alleviate credit constraints), (ii) connect you to someone who is interested in renting out land (revealing an interest in expanding farm size) (iii) connect you to someone who is interested in renting in land (revealing an interest in reducing farm size) (iv) provide you with information on how to improve land quality (revealing an interest in solving land quality issues) (v) put you in touch with someone that can provide reliable labor for your farm (revealing an interest in addressing labor constraints), or (vi) coordinate a meeting with your neighbors to facilitate improvements to the drainage and irrigation system (revealing an interest in addressing water and drainage issues).

Our purpose with this relative ranking analysis was to provide options that were costless to the farmer, yet could tell us about the kinds of problems that seemed solvable to them.

To reduce overburdening with an excessive number of similar questions, we presented farmers with a random selection of 3 of the 15 possible pairwise combinations of these improvement options and asked them to choose their preferred options among the two alternatives presented in each pair. We then aggregated these preferences into a single ranking. The aggregation was done using the Schulze method, which has desirable properties for pairwise comparisons and takes into account the full ranking of preferences rather than just the first

Table 5
Ranking of improvement options in phone survey.

Improvement option	Relative ranking
Provide you with information on how to improve soil quality	1
Coordinate a meeting to facilitate improvements to the drainage and irrigation system	2
Put you in touch with someone for reliable labor	3
Put you in touch with a financial institution	4
Connect you to someone interested in renting out land	5
Connect you to someone interested in renting in land	5

Notes: This table shows the ranking of six improvement options based on responses to 3 randomly selected pairwise questions presented to respondents. The aggregation and ranking of responses is done using the Schulze method, which identifies the strongest option that is preferred to other options when presented together in pairs, and considers as the strongest option the one most preferred by a majority of respondents in pairwise comparisons. Once the first option is selected, the process is repeated to choose the second option among the remaining options in the same manner. The process is then repeated until all options are ranked.

choice(s). The method identifies the strongest option that is preferred to other options when presented together in pairs. The strongest option is the one most preferred by a majority of respondents in pairwise comparisons. Once the first option is selected, the process is repeated to choose the second option. This second option is the one that wins the most in a series of paired match-ups with other options not selected as the first one. The process is then repeated until all options have been ranked.

In Table 5, we present average relative attractiveness of these different options in a decreasing order. The highest-ranked option by farmers was ‘information on how to improve soil quality’. This is in contrast to what we found in Table 2, where land suitability and soil drainage issues were deemed quantitatively unimportant for non-use. We interpret this as something which is of interest to farmers because it could improve outcomes in their continuously used plots (comprising 22.9 percent of their land).

The second most attractive option to farmers was ‘improvements to the irrigation and drainage system’. Strong preference for this option is consistent with the results on non-use among farmers as solution to important causes of non-use in both continuously and intermittently unused land.

7. Conclusion

In this paper, we use satellite imagery to evaluate the impact of irrigation infrastructure development on land cultivation rates in a context where there are no data available on cultivated land and agricultural output over time. We analyze the case of the Senegal River Valley, where a large-scale irrigation development initiative was launched in the 1980s with the construction of two major dams along the Senegal River. Since then, the program has progressively expanded irrigated land through the development of canal infrastructure, ultimately reaching an area of 128,000 hectares.

To fill the land use data gap, we turn to a collection of over 3000 Landsat images collected between 1985 and 2019. We convert the information into a pixel-level indicator of whether land is cultivated or not. We then combine this information with administrative data on irrigation infrastructure projects completed between 1988 and 2019. In doing so, we produce a development unit-level panel of cultivation rates by season spanning over three decades, along with date of project completion and the nature of the project. We use an event-study specification of agricultural land use to identify the effect of irrigation infrastructure.

Our results provide strong evidence that cultivation rates increase markedly beginning in the first year after project completion, especially during the dry season. On average, we find that cultivation rates are remarkably stable at around 25 percentage points above pre-irrigation

levels for the first 20 years, and trend even higher from years 20 to 25. We also show that sensitivity to temperature fluctuations is much diminished with irrigated cultivation, an important finding in the context of climate change.

Analyzing heterogeneity across production units, we find that, despite overall progress in cultivation rates, more than half of production units make either low or intermittent use of the land. High use of the land is seen more frequently in the most recent cohort of projects (2010–19), in projects developed by public agencies rather than privately, and in agribusiness projects. We observe, however, some convergence over time within the three categories of projects.

In spite of the large improvement in cultivation rates since 2000, the non-cultivation problem within completed projects remains a recurrent issue affecting 25 percent of the land as of 2019.

Using a survey of production units, we elicit perceived reasons for under-utilization of irrigation potential. We find that water access problems are a large constraint on cultivation due to deficient infrastructure construction or maintenance issues. Farmers thus consider that improving the irrigation and drainage systems are the most important interventions needed to reduce cultivation failures, and that these interventions require initiatives taken beyond the farm level. These water access constraints are most clearly associated to lands that are continuously unused. In second place, farmers in the Senegal River Valley seem to be constrained by financial or credit constraints, and these are more prevalent among farmers that only used their land intermittently.

More broadly, this research provides methods to exploit publicly available satellite data to conduct ex-post impact evaluation of important policies for developing country agriculture. This research also provides an exciting proof of concept on how satellite-based research can be successfully complemented with on the ground survey data, as a way of validating findings and probing different hypotheses for what is being observed remotely. We hope these tools will be broadly used by other researchers.

One of the most novel findings to come out of these new data sources and methods is the importance of intermittent land usage, a quantitatively important phenomenon that hitherto has not been widely appreciated—most likely because it requires high frequency data on usage patterns over the same plot, something low frequency farmer surveys have trouble identifying. Intermittent land usage is a phenomenon amenable to policy intervention and a potentially low-cost way of boosting agricultural output that may become increasingly relevant in the context of a changing climate.

CRedit authorship contribution statement

Abdoulaye Cisse: Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Alain de Janvry:** Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Joel Ferguson:** Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Marco Gonzalez-Navarro:** Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Samba Mbaye:** Resources, Project administration, Conceptualization. **Elisabeth Sadoulet:** Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Mame Mor Anta Syll:** Project administration, Methodology, Investigation, Formal analysis.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jdevco.2025.103615>.

Data availability

<https://osf.io/2z6n4/>.

References

- Asher, Sam, Campion, Alison, Gollin, Douglas, Novosad, Paul, 2022. The long-run development impacts of agricultural productivity gains: Evidence from irrigation canals in India. CEPR Discussion Paper No. 17414, Centre for Economic Policy Research London, UK.
- BenYishay, Ariel, Sayers, Rachel, Singh, Kunwar, Goodman, Seth, Walker, Madeleine, Traore, Souleymane, Rauschenbach, Mascha, Noltze, Martin, 2024. Irrigation strengthens climate resilience: Long-term evidence from Mali using satellites and surveys. *PNAS Nexus* 3 (2), 1–9.
- Blakeslee, David, Dar, Aaditya, Fishman, Ram, Malik, Samreen, Pellegrina, Heitor S., Bagavathinathan, Karan Singh, 2023. Irrigation and the spatial pattern of local economic development in India. *J. Dev. Econ.* 161, 102997.
- Blakeslee, David, Fishman, Ram, Srinivasan, Veena, 2020. Way down in the hole: Adaptation to long-term water loss in rural India. *Am. Econ. Rev.* 110 (1), 200–224.
- Borusyak, Kirill, Jaravel, Xavier, Spiess, Jann, 2024. Revisiting event study designs: Robust and efficient estimation. *Rev. Econ. Stud.* 91 (6), 3253–3285.
- Bourgoin, Jeremy, Diop, Djibril, 2023. Atlas d'un Territoire en Transition - Regards sur le Delta du Fleuve Sénégal. CIRAD-ISRA.
- Bruckmann, Laurent, 2018. Crue et développement rural dans la vallée du Sénégal: Entre marginalisation et résilience. *Belg. - Rev. Belg. de Géogr.* 2.
- Burke, Marshall, Driscoll, Anne, Lobell, David B., Ermon, Stefano, 2021. Using satellite imagery to understand and promote sustainable development. *Science* 371 (6535).
- Dufo, Esther, Pande, Rohini, 2007. Dams. *Q. J. Econ.* 122 (2), 601–646.
- Faye, M., Fofana, M.B., Bélières, Jean-François, 1995. Présentation de la banque de données de la SAED pour le suivi des aménagements hydro-agricoles et des organisations paysannes. In: Boivin, Pascal, Dia, Ibrahim, Lericollais, André, Poussin, Jean Christophe, Santoir, Christian, Seck, Sidi Mohamed (Eds.), Nianga, Laboratoire de l'Agriculture Irriguée en Moyenne Vallée du Sénégal. In: Colloques et Séminaires, ORSTOM, ISRA, Paris, pp. 513–533.
- Hein, Lars, De Ridder, Nico, 2006. Desertification in the Sahel: A reinterpretation. *Global Change Biol.* 12 (5), 751–758.
- Higginbottom, Thomas P., Adhikari, Roshan, Dimova, Ralitz, Redicker, Sarah, Foster, Timothy, 2021. Performance of large-scale irrigation projects in Sub-Saharan Africa. *Nat. Sustain.* 4 (6), 501–508.
- Higginbottom, Thomas P., Adhikari, Roshan, Foster, Timothy, 2023. Rapid expansion of irrigated agriculture in the Senegal river valley following the 2008 food price crisis. *Environ. Res. Lett.* 18 (1), 014037.
- Hornbeck, Richard, Keskin, Pinar, 2014. The historically evolving impact of the Ogallala aquifer: Agricultural adaptation to groundwater and drought. *Am. Econ. J.: Appl. Econ.* 6 (1), 190–219.
- Jain, Meha, 2020. The benefits and pitfalls of using satellite data for causal inference. *Rev. Environ. Econ. Policy.*
- Le Roy, Xavier, 2000. La difficile mutation de l'agriculture irriguée dans la vallée du fleuve Sénégal. *Territ. Mutat.* 165–177.
- Lobell, David B., Azzari, George, Burke, Marshall, Gourlay, Sydney, Jin, Zhenong, Kilic, Talip, Murray, Siobhan, 2020. Eyes in the sky, boots on the ground: Assessing satellite-and ground-based approaches to crop yield measurement and analysis. *Am. J. Agric. Econ.* 102 (1), 202–219.
- Lobell, David B., Burke, Marshall B., Tebaldi, Claudia, Mastrandrea, Michael D., Falcon, Walter P., Naylor, Rosamond L., 2008. Prioritizing climate change adaptation needs for food security in 2030. *Science* 319 (5863), 607–610.
- Lobell, David B., Di Tommaso, Stefania, You, Calum, Djima, Ismael Yacoubou, Burke, Marshall, Kilic, Talip, 2019. Sight for sorghums: Comparisons of satellite-and ground-based sorghum yield estimates in Mali. *Remote. Sens.* 12 (1), 100.
- MAER, 2015. Rapport de Performance 2014. Ministère de l'Agriculture et de l'Équipement Rural, Direction de l'Analyse, de la Prévision et des Statistiques Agricoles.
- Murillo-Sandoval, Paulo J., Gjerdseth, Emma, Correa-Ayram, Camilo, Wrathall, David, Van Den Hoek, Jamon, Dávalos, Liliana M., Kennedy, Robert, 2021. No peace for the forest: Rapid, widespread land changes in the Andes-Amazon region following the Colombian civil war. *Glob. Environ. Chang.* 69, 102283.
- Myneni, Ranga B., Hall, Forrest G., Sellers, Piers J., Marshak, Alexander L., 1995. The interpretation of spectral vegetation indexes. *IEEE Trans. Geosci. Remote Sens.* 33 (2), 481–486.
- Rosa, Lorenzo, Chiarelli, Davide Danilo, Sangiorgio, Matteo, Beltran-Peña, Areidy Aracely, Rulli, Maria Cristina, D'Odorico, Paolo, Fung, Inez, 2020. Potential for sustainable irrigation expansion in a 3° C warmer climate. *Proc. Natl. Acad. Sci.* 117 (47), 29526–29534.
- Rufin, Philippe, Levers, Christian, Baumann, Matthias, Jägermeyr, Jonas, Krueger, Tobias, Kuemmerle, Tobias, Hostert, Patrick, 2018. Global-scale patterns and determinants of cropping frequency in irrigation dam command areas. *Glob. Environ. Chang.* 50, 110–122.
- SAED, JICA, 2019. Preparatory Survey on Senegal River Valley Irrigated Rice Farming Improvement Project in Republic of Senegal. Japan International Cooperation Agency and Société d'Amenagement et d'Exploitation des Terres du Delta du Fleuve Sénégal et des Vallées du Fleuve Senegal et de la Falémé.
- Strobl, Eric, Strobl, Robert O., 2011. The distributional impact of large dams: Evidence from cropland productivity in Africa. *J. Dev. Econ.* 96 (2), 432–450.
- Sun, Liyang, Abraham, Sarah, 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econometrics* 225 (2), 175–199.
- Tattheo, Mokganedi, D'Odorico, Paolo, 2022. Are African irrigation dam projects for large-scale agribusiness or small-scale farmers? *Environ. Res. Commun.* 4 (1), 432–450.
- UN, 2024. World Population Prospects 2024: Summary of Results. United Nations Department of Economic and Social Affairs, Population Division.
- USAID, 2017. Évaluation des Expériences et du Potentiel Hydroagricole pour la Petite Irrigation dans les Zones Nord et Sud du Sénégal. Feed the Future Senegal-Naatal Mbay, U.S. Agency for International Development.
- Wang, Di, Zhang, Peng, Chen, Shuai, Zhang, Ning, 2024. Adaptation to temperature extremes in Chinese agriculture, 1981 to 2010. *J. Dev. Econ.* 166, 103196.
- Welch, Jarrod R., Vincent, Jeffrey R., Auffhammer, Maximilian, Moya, Piedad F., Dobermann, Achim, Dawe, David, 2010. Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures. *Proc. Natl. Acad. Sci.* 107 (33), 14562–14567.
- Wuepper, David, Oluoch, Wyclife Agumba, Hadi, Hadi, 2025. Satellite data in agricultural and environmental economics: Theory and practice. *Agric. Econ.* 56 (3), 493–511.
- Wuepper, David, Wang, Haoyu, Schlenker, Wolfram, Jain, Meha, Finger, Robert, 2023. Institutions and Global Crop Yields. NBER Working Paper 31426.
- Xue, Jinru, Su, Baofeng, et al., 2017. Significant remote sensing vegetation indices: A review of developments and applications. *J. Sens.* 2017.
- Zwart, Sander J., 2017. A detailed map of irrigation infrastructure along the Senegal River by 2015. <http://dx.doi.org/10.13140/RG.2.2.28009.49761>, AfricaRice Center GIS Report 4.